

# Semi-supervised Discovery of Informative Tweets During the Emerging Disasters

Shanshan Zhang  
Temple University  
1805 N. Broad Street  
Philadelphia, PA

zhang.shanshan@temple.edu

Slobodan Vucetic  
Temple University  
1805 N. Broad Street  
Philadelphia, PA  
vucetic@temple.edu

## ABSTRACT

The first objective towards the effective use of microblogging services such as Twitter for situational awareness during the emerging disasters is discovery of the disaster-related postings. Given the wide range of possible disasters, using a pre-selected set of disaster-related keywords for the discovery is suboptimal. An alternative that we focus on in this work is to train a classifier using a small set of labeled postings that are becoming available as a disaster is emerging. Our hypothesis is that utilizing large quantities of historical microblogs could improve the quality of classification, as compared to training a classifier only on the labeled data. We propose to use unlabeled microblogs to cluster words into a limited number of clusters and use the word clusters as features for classification. To evaluate the proposed semi-supervised approach, we used Twitter data from 6 different disasters. Our results indicate that when the number of labeled tweets is 100 or less, the proposed approach is superior to the standard classification based on the bag of words feature representation. Our results also reveal that the choice of the unlabeled corpus, the choice of word clustering algorithm, and the choice of hyperparameters can have a significant impact on the classification accuracy.

## CCS Concepts

•Information systems → Data analytics;

## Keywords

semi-supervised learning; Twitter; crisis management

## 1. INTRODUCTION

During the past decade, there has been plentiful evidence that social media can be very useful in improving the situational awareness during emergencies and disasters [4], and several research groups even proposed specific systems to extract, process, and summarize social media data for such purpose [1, 2, 5]. The prototypical situational awareness system relying on social media, such as the popular microblogging service Twitter, consists of a data collection, a data processing, and a visual interface component. The first critical challenge in making those systems work is selection of the disaster-related posts from a massive stream of posts pouring from services such as Twitter. The objective of this work is to propose a new approach for selection

of informative posts, with a particular focus on the beginning stages of a disaster, when very little is known about the emergency and the ways in which people report about and react to it.

There are two main approaches for selection of the informative posts. One is the information retrieval approach, which is based on developing a lexicon of disaster-related unigrams or bigrams based on a study of posts from previous disasters and emergencies [8, 10]. While there are terms that are commonly used over different disaster types (e.g. “victims”, “prayers”, “help”), each disaster is different and might require its own specialized lexicon. As such, pre-built lexicons cannot be expected to have very large sensitivity and specificity during the emerging disasters. An alternative supervised learning approach is based on training a classifier that recognizes informative posts [1, 2, 5]. To train the classifier, there is a need to have access to labeled posts. If a sufficient number of labeled posts are available, the supervised approach is superior to the keyword-based approach. However, during the first few minutes or hours from the emergence of a disaster, the number of labeled posts might be too small to result in accurate classification.

The main issue when training from small labeled data is avoiding to build overly complex models that cannot generalize well. A typical remedy is to keep classification models simple. That is why logistic regression is a very popular model choice in the low-sample learning scenarios. A closely related issue is that the dimensionality of data is often too large and it complicates the issue of learning from small labeled data. For example, in case of classifying text data such as microblogs, a popular data representation approach is the bag of words representation, in which each word from a dictionary is treated as a feature. A standard remedy is to reduce dimensionality by feature selection and use some form of model regularization. However, when the number of labels is very small (less than a hundred), selection of the most informative words to be used as features can become unreliable. In particular, for all but the most frequent words, there might be too little evidence to estimate how useful they are with any degree of statistical certainty.

In this work, we propose to cluster similar words in groups and treat each cluster as a feature. By treating word clusters as features, the features would become less sparse and the estimate of their usefulness for classification would be more confident. In addition, word clusters would allow utilizing words that occur rarely or do not even occur in the labeled corpus. To create word clusters, we rely on unlabeled historical posts. One of the open questions in this study was

what kind of unlabeled corpus is the most appropriate for this application. In the following text, we describe the details of our approach and provide an experimental evaluation of the approach on Twitter data from several disasters.

## 2. METHODOLOGY

### 2.1 Problem Setup

Let us suppose we are given a corpus of  $n$  labeled documents  $\mathbf{D}_{train} = \{(d_1, y_1), (d_2, y_2), \dots, (d_n, y_n)\}$  and each document  $d_i$  is converted into a feature vector  $\mathbf{x}_i \in \mathbb{R}^m$ . For example, we could use the bag of words representation, where given a vocabulary  $\mathbf{V}$  the feature vector  $\mathbf{x}_i$  is a binary vector of length  $|\mathbf{V}|$ , whose  $j$ -th element indicates whether the  $j$ -th word from the vocabulary is present in the document  $d_i$ . By a slight extension,  $n$ -grams, part-of-speech tags, named entities, or word clusters could also be used for feature representation of documents. The label  $y_i$  of document  $d_i$  is a binary variable indicating whether it is disaster-related or not, where  $y_i = 1$  indicates that the document is disaster related and  $y_i = 0$  that it is not. Let us also assume that we are given a corpus of unlabeled documents  $\mathbf{D}_{unl} = \{d_1, d_2, \dots, d_N\}$  containing  $N$  documents that are available at the training time. The objective is to train a model  $f(x)$  whose output can be used to classify a document. We will assume that the output is real-valued and that larger values indicate a stronger likelihood of being disaster-related. For classification, we need to select a particular threshold  $\theta$  and classify all documents with outputs above the threshold as disaster-related and those below the threshold as unrelated.

### 2.2 Feature Selection Filters

The objective of feature selection is to reduce dimensionality of the feature vector. Feature selection filters are the simplest class of such algorithms that attempt to select only a subset of the most informative features and remove the rest. They do it so by calculating the score of each feature that measures how strong its correlation is with the classification label. Only the  $K$  features with the largest score are retained. There are many known scoring functions [9], among which the  $\chi^2$  statistics and Pointwise Mutual Information (PMI) are very popular for text classification. In this paper, we will use the PMI score, which is defined in the following way:

$$score(t) = PMI(t, pos) - PMI(t, neg) = \log_2 \frac{p(t|pos)}{p(t|neg)} \quad (1)$$

where  $p(t|pos)$  is the probability of feature  $t$  appearing in positive class, and  $p(t|neg)$  is the probability of feature  $t$  appearing in negative class.

### 2.3 Word Clusters as Features

The idea of document representation with word clusters is to group similar words and treat each cluster as a feature. In particular, given a vocabulary  $\mathbf{V}$  the objective is to assign them to one of  $K$  clusters. Given those clusters, each document  $d_i$  is converted into a binary feature vector  $\mathbf{x}_i$  of length  $K$  whose  $j$ -th element  $x_{ij}$  equals one if any of the words from the  $j$ -th cluster are present in the document and zero if not.

To create the clusters, we rely on the corpus of unlabeled documents  $\mathbf{D}_{unl} = \{d_1, d_2, \dots, d_N\}$ . Given the corpus,

we will evaluate two ways of generating word clusters, as explained next.

**Brown clustering (BC)**[3]. This is a traditional algorithm for word clustering. It is a hierarchical clustering algorithm that assigns a single cluster to each word and proceeds by merging the two clusters that result in the smallest reduction of global mutual information. The output of Brown clustering is a dendrogram of words, which means we can cut at any level in the hierarchy to create word clusters.

**Clustering based on the skip-gram representation of words.** The idea of neural language models, of which the skip-gram model [7] is a representative, is to learn low-dimensional vector representation of words. The skip-gram model defines the probability of word  $w_N$  being in the neighbourhood of a target word  $w_T$  in text by assuming that the words in the neighborhood are conditionally independent given the target word. It models probability of neighbour  $w_N$  given the target word  $w_T$  as a softmax function

$$p(w_N|w_T) = \frac{\exp(v_{w_N}'^T v_{w_T})}{\sum_{w=1}^{|\mathbf{V}|} \exp(v_w'^T v_{w_T})} \quad (2)$$

where  $v_w \in \mathbb{R}^p$  and  $v_w' \in \mathbb{R}^p$  are the “input” and “output”  $p$ -dimensional vector representations of word  $w$ . The input and output representations of every word are obtained by maximizing the log-likelihood of the model over a corpus of documents. In this paper, we use word2vec [7], which is a stochastic gradient algorithm that is commonly used to maximize this objective function. The resulting word representations place words occurring in similar contexts near each other.

Given the output word representations from the word2vec algorithm, we use the traditional k-means clustering algorithm to group all words from the vocabulary into  $K$  clusters.

## 3. EXPERIMENTS AND EVALUATION

The objective of our experimental evaluation is to obtain an insight into the performance of the proposed semi-supervised approach, which uses an unlabeled document corpus to create word clusters and trains a classifier on a small labeled document corpus where the word clusters are used as features.

### 3.1 Data Set

For our experiments, we used CrisisLexT6 data described in [8]. It is a collection of tweets related to 6 natural disasters that occurred between October 2012 and July 2013: Sandy Hurricane (**SH**), Boston Bombings (**BB**), Oklahoma Tornado (**OT**), West Texas Explosion (**WE**), Alberta Floods (**AF**), Queensland Floods (**QF**).

For each disaster, the same number of tweets was selected using two different methods. One method relied on a predefined set of keywords appropriate for a particular disaster. Another method relies on selecting geocoded tweets originating from the affected areas during the onset of a disaster. The collected tweets were labeled as disaster or not disaster related by 100 workers from Crowdfunder. Out of the 60,000 Twitter IDs of the labeled tweets provided by the authors of [8], we were only able to download 70% of them and we summarize their basic statistics in Table 1.

From Table 1, we observe a highly skewed distribution of classes depending on the selection criteria. In particu-

Table 1: Basic data statistics of each disaster.

Disaster Type	# Location-based sample	# Positive	# Keyword-based sample	# Positive
Sandy Hurricane	3505	1000	3267	3058
Boston Bombings	3604	610	3651	3451
Oklahoma Tornado	4088	423	3897	3362
West Texas Explosion	4535	305	4245	4168
Alberta Floods	3712	367	4384	4044
Queensland Floods	3851	361	4302	4051

lar, the keyword-based sample contains a high fraction of positive (disaster-related) labels, while the location-based sample has a high fraction of negative labels. Unlike the previous work that merged tweets from those two samples and split them randomly into training and test sets for evaluation [5, 6, 8], in this work we used only the tweets obtained by the location search. We discarded the portion of tweets obtained by the keyword-based selection. Our reasoning was that as long as the keywords are well defined by authorities, virtually all tweets sampled by the keyword search can be assumed as positives. Identifying disaster-related Tweets within the location-based sample is a more challenging classification problem, which can be particularly important during the initial minutes or hours of an emerging disaster.

## 3.2 Experiment Design

In our experiments we wanted to compare the proposed approach with the standard supervised learning approach that uses feature selection on the bag of words and trains a logistic regression classifier. We also wanted to explore impact of a range of choices on the accuracy, including the choice of the word clustering algorithm, the choice of an unlabeled document corpus, the size of the labeled corpus, the number of features, and the choice of hyperparameters.

*Preprocessing.* All tweets were tokenized using the tool from `twitter_nlp`<sup>1</sup>. Stop-words, URLs, and user mentions (`@username`), tokens too short (2 characters or less), tokens too long (16 characters or more) were removed from the tweets. Given the tweets preprocessed in this way, we divided the location-based tweets of each disaster randomly into two subsets: 70% as training set and 30% as the test set. The test set was not used in any stage of classifier building and was used solely to calculate the accuracy.

*Classification Algorithm.* In all experiments we relied on Logistic Regression (**LR**) with  $L_2$  regularization. The regularization hyperparameter was obtained using leave-one-out cross-validation on the training data.

*Choice of training size.* We varied the training size in the range  $t = [20, 50, 100, 200, 500, 1000]$ . This allowed to achieve an insight into the impact of the training size on accuracy and its relative impact on different algorithms. The smallest training sizes of 20 and 50 are representative of scenarios with extremely limited labeled data sets that might be expected shortly after the outset of a disaster, while the larger training sizes are representatives of labeled data sets that can become available with a more significant delay.

*Choice of features.* We compared two clustering algorithms described in the Methodology, the Brown Clustering (**BC**) and the clustering based on word2vec (**W2V**). We denote feature selection based on the traditional bag-of-words

as **BOW**. We varied the number of clusters in the range of  $K = [50, 100, 200, 500, 1000, 2000]$ . To be comparable, for **BOW** we used feature selection based on PMI score to train Logistic Regression based on the top  $K$  scoring words.

*Choice of the unlabeled corpus for word clustering.* We explored several types of unlabeled Twitter sets for word clustering using **BC** and **W2V** approaches:

- **D**<sub>50K</sub>. For any disaster, the unlabeled corpus contains all available CrisisLexT6 tweets from the other 5 disasters.
- **D**<sub>50K<sub>p</sub></sub>. For any disaster, the unlabeled corpus contains all positively labeled CrisisLexT6 tweets from the other 5 disasters.
- **D**<sub>3K</sub>. For any disaster the corpus contains all available keyword-based tweets from that disaster.

In addition, we downloaded a pre-trained *word2vec* matrix **W** with word vector dimension  $p = 400$  and vocabulary size  $|\mathbf{V}| \sim 3,000,000$  that was trained on 400 million Tweets.

By combining different choices for feature construction, we obtained 8 different algorithms that we evaluated experimentally: bag of words (**BOW**), bag of Brown clusters (**BC**<sub>50K</sub>, **BC**<sub>50K<sub>p</sub></sub>, **BC**<sub>3K</sub>), and bag of word2vec clusters (**W2V**<sub>50K</sub>, **W2V**<sub>50K<sub>p</sub></sub>, **W2V**<sub>3K</sub>, **W2V**<sub>400M</sub>)

*Other experimental settings.* We run *word2vec* with negative sampling rate  $10^{-3}$ , context window size 100, and word vector dimension 20 (due to a relatively small unlabeled corpus). For each choice of the feature construction algorithms, each training set size, and each number of features  $K$ , we run Logistic Regression 10 times on randomly sampled subsets from the training set. The average Area Under the ROC Curve (AUC) of the 10 random sampling experiments is reported.

## 3.3 Results

Our results provide insights into several interesting questions related to recognizing disaster-related tweets during the emerging disasters. Figure 1 is the most comprehensive view into our results. For each disaster and each training size, we are showing 8 curves, representing 8 different feature construction algorithms, as a function of the number of features. However, to answer our posed questions, we summarized those results with 4 tables, as will be described in the following.

**Are word clusters better than BOW?** In Table 2, we compare the accuracies of the best word clustering and the best **BOW** classifier for each of the 6 disasters and for each training data size. Notice that for **BOW** approach at a specific training size, we report the best accuracy among different choices of  $K$ , which means we are treating the  $K$  as a hyperparameter during training. For the 7 semi-supervised approaches, we report the best accuracy from that group. When the training data size is 20, 50, and 100, the word clustering approach is superior to **BOW** on all but the Oklahoma Tornado (**OT**) data set. This result strongly indicates that an unlabeled Twitter corpus could provide a very useful information when there is a severe deficit of labeled tweets for an emerging disaster. As the number of labeled tweets is exceeding 100, the benefit from the unlabeled corpus fades, to the extent that **BOW** is more accurate than the word clusters on all 6 disasters when the training data size reaches 1,000. This is an expected result, since such a large training data set is sufficient to identify the most informative disaster-related words.

<sup>1</sup>[https://github.com/aritter/twitter\\_nlp](https://github.com/aritter/twitter_nlp)

Table 2: Summary results to compare Semi-supervised (SS) and BOW algorithms.

Train	SH		BB		OT	
	SS*	BOW*	SS*	BOW*	SS*	BOW*
20	<b>0.734</b>	0.69	<b>0.714</b>	0.7	0.688	0.693
50	<b>0.76</b>	<b>0.76</b>	<b>0.761</b>	0.707	0.746	0.746
100	<b>0.8</b>	<b>0.8</b>	<b>0.794</b>	0.774	0.795	0.803
200	0.838	0.844	<b>0.816</b>	0.811	0.846	0.857
500	0.865	0.878	0.844	0.874	0.887	0.897
1000	0.883	0.892	0.868	0.892	0.909	0.918

Train	WE		AF		QF	
	SS*	BOW*	SS*	BOW*	SS*	BOW*
20	<b>0.754</b>	0.709	<b>0.746</b>	0.686	<b>0.716</b>	0.645
50	<b>0.77</b>	0.774	<b>0.798</b>	0.753	<b>0.764</b>	0.702
100	<b>0.83</b>	0.807	<b>0.851</b>	0.848	<b>0.812</b>	0.783
200	0.863	0.865	<b>0.907</b>	0.876	<b>0.856</b>	0.85
500	0.93	0.921	<b>0.936</b>	0.928	<b>0.9</b>	0.9
1000	0.935	0.947	0.942	0.949	0.929	0.934

Does the choice of a word clustering algorithm make a difference? In Table 3 we compare the best accuracy obtained by BC clustering to the best accuracy by W2V clustering. For both BC and W2V, we only considered  $D_{50K}$ ,  $D_{50Kp}$ , and  $D_{3K}$  corpuses. None of the algorithms is superior over all 6 disasters. W2V is superior on OT, AF, and QF, BC is superior on SH, and they are similar on the remaining two disasters. This is a slightly surprising result, considering the age of Brown clustering and the recent hype surrounding the word2vec algorithm. Studying this result more in depth would certainly be worthwhile. For the purposes of this paper, our conclusion is that both algorithms have their strengths, but that if we have to choose one, it would be W2V.

Table 3: Comparing two clustering algorithms BC and W2V.

Train	SH		BB		OT	
	BC	W2V	BC	W2V	BC	W2V
20	<b>0.734</b>	0.65	0.707	<b>0.714</b>	0.639	<b>0.682</b>
50	<b>0.758</b>	0.736	<b>0.761</b>	0.755	0.706	<b>0.741</b>
100	<b>0.792</b>	0.764	<b>0.794</b>	0.784	0.762	<b>0.789</b>

Train	WE		AF		QF	
	BC	W2V	BC	W2V	BC	W2V
20	0.73	<b>0.754</b>	0.575	<b>0.613</b>	0.591	<b>0.651</b>
50	<b>0.759</b>	0.748	0.617	<b>0.678</b>	0.668	<b>0.702</b>
100	0.801	<b>0.82</b>	0.681	<b>0.689</b>	0.724	<b>0.761</b>

#### What is the impact of the unlabeled corpus choice?

In Table 4 we are showing the accuracies of the best W2V algorithm for four different choices of an unlabeled corpus:  $D_{50K}$ ,  $D_{50Kp}$ ,  $D_{3K}$ , and  $D_{400M}$ . On four disasters (SH, OT, AF, QF), the 400 million Twitter corpus results in the best accuracy, it is competitive on WE, and only on BB it is less accurate than  $D_{50K}$ . This result is a strong indicator that mantra “the more data the better” works in the emerging disaster scenario. Confirming this, it is interesting to see that using  $D_{50K}$  is in general better than using its positively labeled subset  $D_{50Kp}$ . Finally, using the location-based tweets from the emerging disaster itself does not seem to be particularly effective, and it could probably be attributed to the small size of this corpus.

How many word clusters? In Table 5 we show the best choice of the number of clusters for W2V obtained from the 400 million corpus. When the training data size is small, the smaller number of clusters is preferable. This result illustrates that the benefit of merging similar words into a

Table 4: Comparing 4 types of unlabeled corpus.

Train	SH				BB			
	3K	50Kp	50K	400M	3K	50Kp	50K	400M
20	0.646	0.693	0.665	<b>0.696</b>	0.694	0.658	<b>0.714</b>	0.662
50	0.681	0.726	0.736	<b>0.746</b>	0.722	0.683	<b>0.755</b>	0.737
100	0.661	0.776	0.764	<b>0.795</b>	0.737	0.733	<b>0.784</b>	0.768

Train	OT				WE			
	3K	50Kp	50K	400M	3K	50Kp	50K	400M
20	0.604	0.651	0.682	<b>0.688</b>	0.661	0.702	<b>0.754</b>	0.71
50	0.692	0.708	0.741	<b>0.746</b>	0.719	0.702	0.748	<b>0.752</b>
100	0.695	0.732	0.789	<b>0.795</b>	0.774	0.777	<b>0.82</b>	<b>0.82</b>

Train	AF				QF			
	3K	50Kp	50K	400M	3K	50Kp	50K	400M
20	0.725	0.608	0.613	<b>0.746</b>	0.608	0.589	0.651	<b>0.716</b>
50	0.781	0.636	0.678	<b>0.798</b>	0.636	0.653	0.702	<b>0.764</b>
100	0.811	0.669	0.689	<b>0.851</b>	0.703	0.698	0.761	<b>0.812</b>

smaller number of clusters to create dense features outweighs a potential loss of semantic cohesion in those clusters. As expected, for larger training sets, cluster cohesion becomes an important driving factor for high accuracy.

Table 5: The number of clusters (K) to achieve best performance for W2V<sub>400M</sub>

Train	SH	BB	OT	WE	AF	QF
20	500	2000	500	100	500	200
50	1000	1000	100	1000	1000	100
100	1000	2000	500	500	2000	1000
200	2000	2000	500	2000	2000	2000
500	2000	2000	2000	2000	1000	1000
1000	2000	2000	2000	1000	2000	1000

Impact of training size. Not to be lost in the previous discussion, it is evident that the number of labeled tweets has a strong influence on accuracy, and we can observe a robust increase in accuracy with the training data size.

## 4. CONCLUSION

In this paper we addressed the problem of retrieving disaster-related tweets shortly after the onset of a disaster. In such a scenario, we can expect to have access to a very limited number of labeled tweets. The accuracy of classifiers trained on such small data would be limited. To remedy this problem, we proposed a semi-supervised approach that can utilize a large unlabeled corpus of tweets to create word clusters and use them as features for classification. Our experiments on Twitter data from 6 disasters strongly indicate that the proposed semi-supervised approach could most often result in the improvements in accuracy as compared to the traditional supervised learning approach that uses feature selection on the bag of words features. Our study also provides useful insights into different modeling choices when using the proposed approach. While “the bigger the better” mantra in data science mostly holds true in this application, by a careful look at the results, it is also evident that “one size does not fit all,” and that many modelling choices have differing effect on different types of disasters. More effort is needed to gain a better insight into those differences and design approaches with more consistent behavior across a wide range of situational awareness scenarios.

## ACKNOWLEDGEMENTS

This work was supported by NSF grant CNS-1461932.

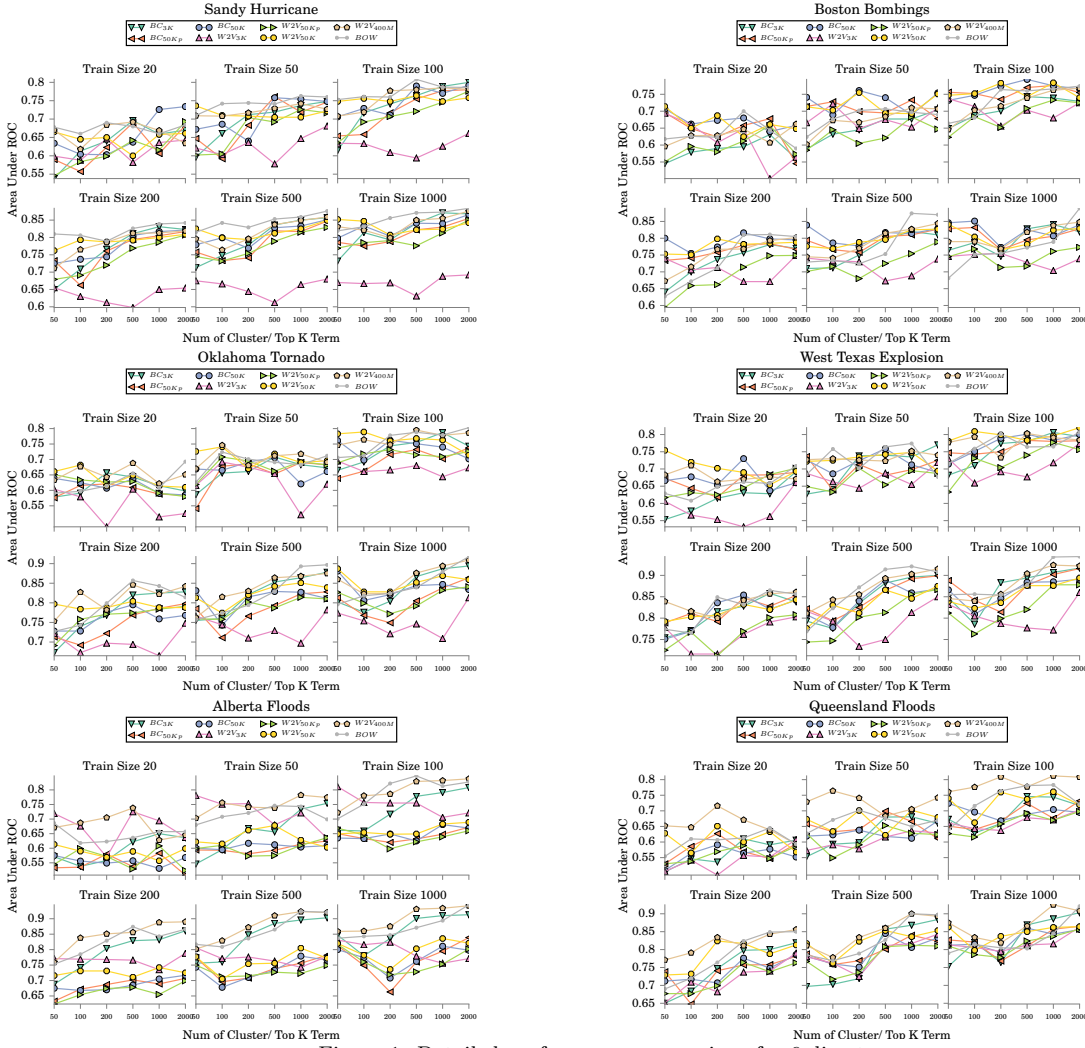


Figure 1: Detailed performance comparison for 6 disasters.

## References

- [1] F. Abel, C. Hauff, G.-J. Houben, R. Stronkman, and K. Tao. Twitcident: fighting fire with information from social web streams. In *Proceedings of the 21st International Conference on World Wide Web*, pages 305–308. ACM, 2012.
- [2] Z. Ashktorab, C. Brown, M. Nandi, and A. Culotta. Tweedr: Mining twitter to inform disaster response. *Proc. of ISCRAM*, 2014.
- [3] P. F. Brown, P. V. Desouza, R. L. Mercer, V. J. D. Pietra, and J. C. Lai. Class-based n-gram models of natural language. *Computational linguistics*, 18(4):467–479, 1992.
- [4] M. Imran, C. Castillo, F. Diaz, and S. Vieweg. Processing social media messages in mass emergency: A survey. *ACM Computing Surveys (CSUR)*, 47(4):67, 2015.
- [5] M. Imran, C. Castillo, J. Lucas, P. Meier, and S. Vieweg. Aidr: Artificial intelligence for disaster response. In *Proceedings of the 23rd International Conference on World Wide Web*, pages 159–162. ACM, 2014.
- [6] M. Imran, P. Mitra, and J. Srivastava. Cross-language domain adaptation for classifying crisis-related short messages. *arXiv preprint arXiv:1602.05388*, 2016.
- [7] T. Mikolov and J. Dean. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 2013.
- [8] A. Olteanu, C. Castillo, F. Diaz, and S. Vieweg. Crisislex: A lexicon for collecting and filtering microblogged communications in crises. In *ICWSM*, 2014.
- [9] P. Radivojac, Z. Obradovic, A. Dunker, and S. Vucetic. Characterization of permutation tests for feature selection. In *Proc. 15th European Conference on Machine Learning*, pages 334–346, 2004.
- [10] I. Temnikova, C. Castillo, and S. Vieweg. Emterms 1.0: a terminological resource for crisis tweets. In *ISCRAM 2015 proceedings of the 12th international conference on information systems for crisis response and management*, 2015.